A More Samples Generated by TinyGAN

Here we show more samples and interpolation results generated by the proposed TinyGAN, as well as some samples by SNGAN-projection for comparison.

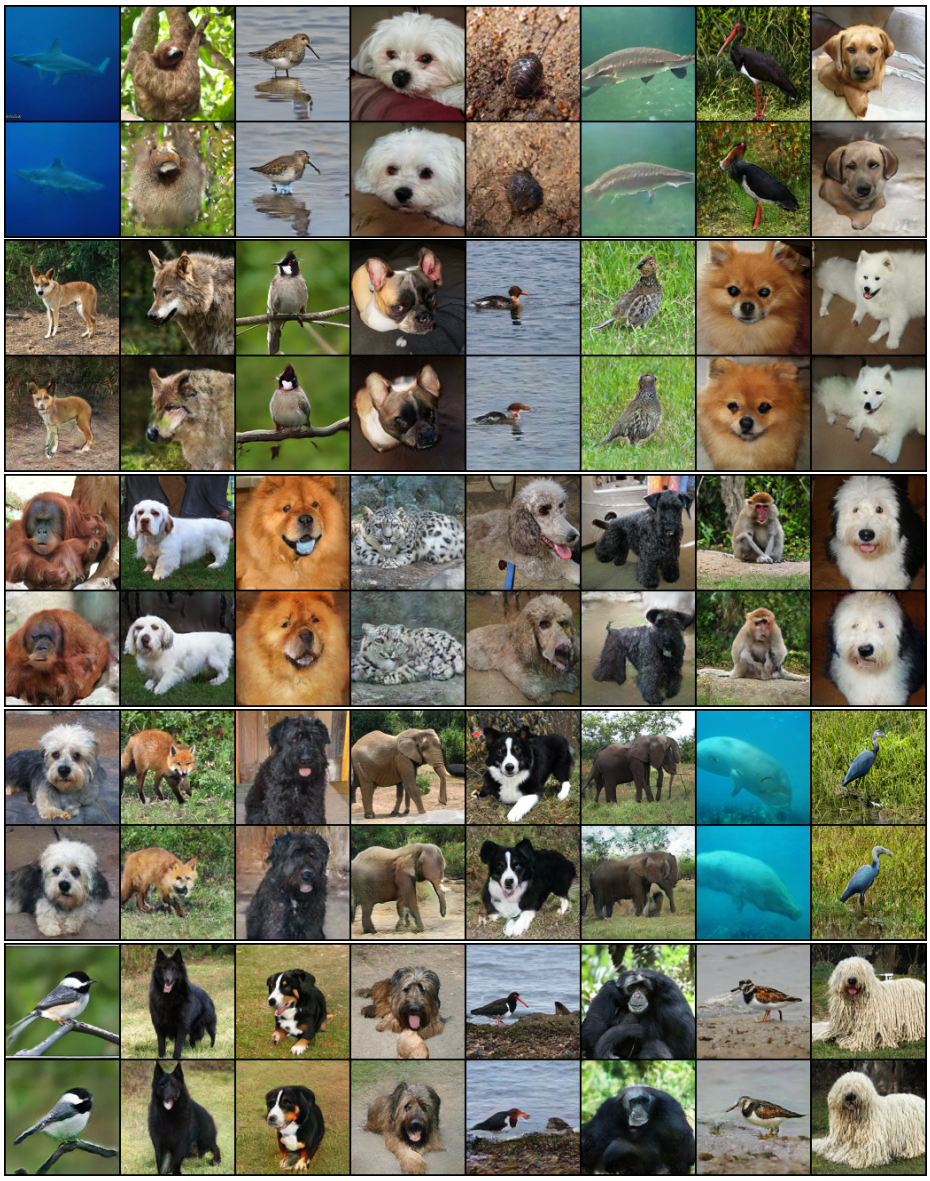

Fig. 12: Images in the odd rows are produced by BigGAN; the even rows are generated by TinyGAN-dw given the same input.

Fig. 13: Randomly sampled images from SNGAN-projection.

Fig. 14: Randomly sampled images from TinyGAN-dw.

Fig. 15: Interpolations between z, y pairs. The first row shows two end points generated by TinyGAN-dw. The second and third rows interpolate between y with z held constant, while the last two rows interpolate between z with y unchanged.

B Results on All Classes

While we are able to achieve satisfactory results when focusing on more homogeneous classes, such as all animal classes, we have less success when training TinyGAN on all 1000 classes of ImageNet. We found it hard for TinyGAN to model objects in very different categories simultaneously, especially for objects with complex textures. Figure 16 shows some failure examples with blurry or unrealistic patches.

A simple solution may be to divide the classes into a small number of groups (e.g. based on ImageNet's ontology), and train a separate TinyGAN model for each group, so that each model only needs to handle a group of more homogeneous classes. Another possible solution may be to use a more sophisticated training schedule and perhaps also more careful hyperparameter-tuning. As the focus of our work is to demonstrate the possibility of obtaining a small-size GAN with an efficient and stable training process, we leave the task of training a single small-size GAN for all the 1000 classes as our future work.

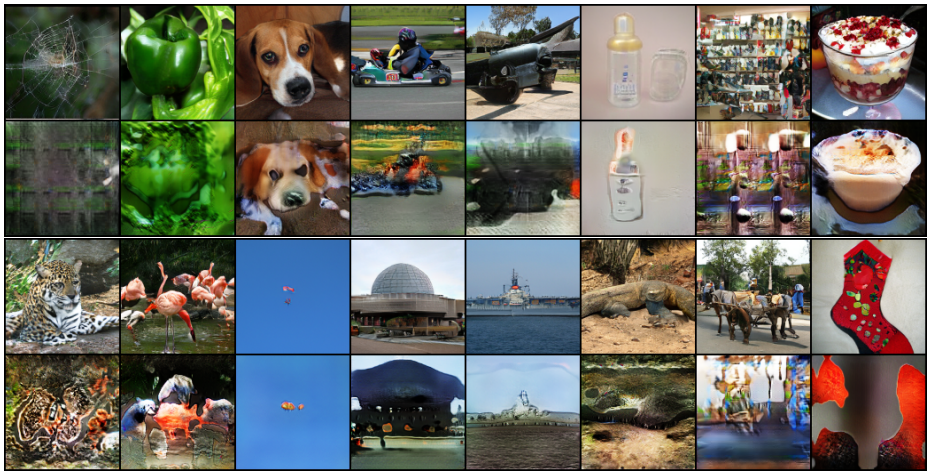

Fig. 16: TinyGAN's failure on modeling all 1000 classes in ImageNet simultaneously. Images in the odd rows are produced by BigGAN, while those in the even rows are by TinyGAN given the same input.